arXiv:1411.3904v2 [stat.ME] 23 Dec 2014

AUTOCORRELATION TYPE FUNCTIONS FOR BIG AND DIRTY DATA SERIES

CHRISTOPH BANDT

One form of big data are signals - time series of consecutive values. In physical experiments, billions of values can now be measured within a second [1, 2]. Signals of heart and brain in intensive care, as well as seismic waves, are measured with 100 up to 1000 Hz over hours, days or even years [3, 4]. A song of 3 minutes on CD comprises 16 million values.

Music and seismic vibrations basically consist of harmonic oscillations for which classical tools like autocorrelation and spectrogram work well [5, 6, 7]. This note presents similar tools for all kinds of rhythmic processes, with non-linear distortion, artefacts, and outliers. Permutation entropy [8, 9, 10] has been used in physics [2, 11], medicine [12, 13, 14], and engineering [15, 16]. Now ordinal patterns [17, 18, 19] are studied in detail for big data. As new version of permutation entropy, we define a distance to white noise consisting of four curious components. Applications to a variety of medical and sensor data are discussed.

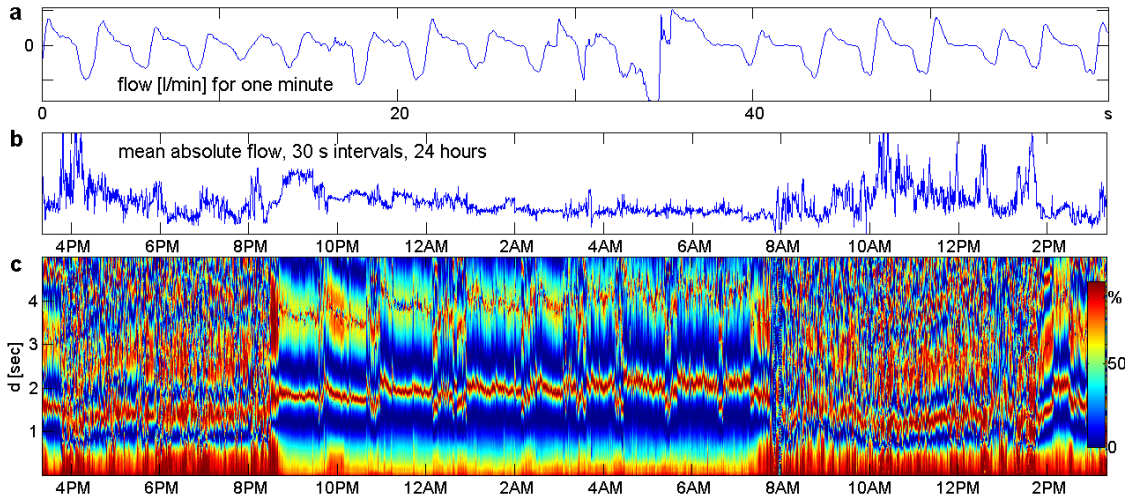


FIGURE 1. Nose breathing of a healthy volunteer in normal life. a: One minute of data without artefacts. b: Mean absolute flow for each 30 seconds of 24 hours measurement. c: New function $\tilde{\tau}(d)$ for each minute shows more structure, needs no calibration of data.

Date: December 24, 2014.

Classical time series from quarterly reports of companies, monthly unemployment figures, daily statistics of accidents etc. consist of 20 up to a few thousand values which were determined with great care. Machine-generated data amount to millions and are not double-checked. Usually it is easy to let the machine run faster and longer but processing goes the other way: data are smoothed out by downsampling. This note shows how high time resolution can be exploited with simple ordinal autocorrelation functions. To demonstrate the remarkable effect of these tools, they will be applied to raw data: no preprocessing, no filters, no detection of outliers, no removal of artefacts, and no polishing of the results.

Figure 1 comes from ongoing work with Achim Beule (Department ENT, Head and Neck Surgery, University of Greifswald) on respiration of healthy volunteers in everyday life. Sensors measuring air flow with sampling frequency 50 Hz were fixed in both nostrils, gently enough to ensure comfort for 24 hours measurement. Mouth breathing was not controlled, so the signals contain lots of artefacts. Traditional analysis takes averages over 30 seconds to obtain a better signal. More information is found in a function $\tilde{\tau}(d)$ for each minute of the dirty signal. As explained below, $\tilde{\tau}$ measures the percentage of "elliptical symmetry" of the signal, depending on time and the delay d which runs from 0.02 to 5 seconds. The collection of these functions, visualized like a spectrogram, shows phases of activity and sleep, various interruptions of sleep, inaccurate measurements around 8 am, a little nap after 2 pm. Frequency of respiration can be read from the lower dark red stripe which marks half of the wavelength, and the upper red and yellow stripe marks the full wavelength (4 seconds in sleep, less than 3 in daily life).

1. Up-down balance. Only the order relation between the values of the time series $x = (x_1, ..., x_T)$ will be used, not the values themselves. In the simplest case, the question is whether the time series goes more often upwards than downwards. For a delay d between 1 and $T/2$, let $n_{12}(d)$ and $n_{21}(d)$ denote the number of time points t for which $x_t < x_{t+d}$ and $x_t > x_{t+d}$, respectively. We determine the relative frequencies of increase and decrease over d steps: $p_{12}(d) = n_{12}(d)/(n_{12}(d) + n_{21}(d))$ and $p_{21}(d) = 1 - p_{12}(d)$. Ties are disregarded. The up-down balance

$$\beta(d) = p_{12}(d) - p_{21}(d) \quad \text{for} \quad d = 1, 2, ...$$

is a kind of autocorrelation function. It reflects the dependence structure of the underlying process and has nothing to do with the size of the data.

For illustration we take water levels from the database of the National Ocean Service [20]. Since water tends to come fast and disappear more slowly, we could expect $\beta(d)$ to be negative, at least for small d. This is mostly true for water levels from lakes, like Milwaukee in Figure 2. For sea stations, tidal influence makes the data almost periodic, with a period slightly smaller than 25 hours. The data have 6 minute intervals and we measure d in hours, writing $d = 250$ as $d = 25$h. A strictly periodic time series with period L fulfils $\beta(L-d) = -\beta(d)$ which in the case $L = 25$ implies $\beta(12.5) = 0$, visible at all sea stations in Figure 2. Otherwise there are big differences: at Honolulu and Baltimore the water level is more likely to fall within the next few hours, at San Francisco it is more likely to increase, and at Anchorage

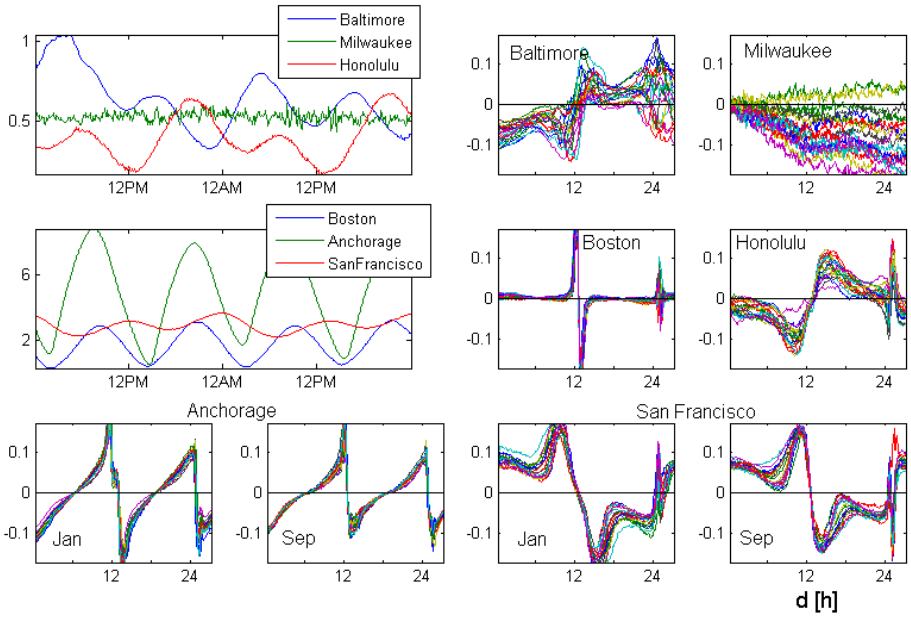


FIGURE 2. Water levels at 6 minute intervals from [20]. Original data shown for Sep 1-2, 2014. Functions β are given for September of 18 years 1997-2014, and also for January in case of Anchorage and San Francisco. d runs from 6 min to 27 hours.

there is a change at 6 hours. Each station has its specific β-profile, almost unchanged during 18 years, which characterizes its coastal shape. β can also change with the season, but these differences are smaller than those between stations.

Figure 2 indicates that β, as well as related functions below, can solve some basic problems of statistics: describe, distinguish, and classify objects. Cf. Figure 13.

2. Sliding windows and statistical variation. The time series is divided into pieces, so-called windows, and β is determined for every window. If we have some hundred windows, which may overlap, we can compose them in a map like Figure 1 or 4, where each time point corresponds to a window, the β-functions are represented by vertical lines, and the values $\beta(d)$ are color coded. Such maps show whether the regime of the underlying process changes in time, and indicate transitions between different regimes: activity and rest, or sleep stages.

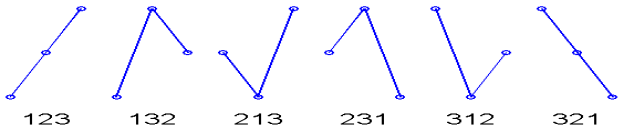
If there are no systematic changes in the appearance of the functions, we can consider the underlying process as stationary and estimate statistical fluctuations. Mathematical arguments say that the standard deviation of $\beta(d)$ for fixed d is $\frac{c}{\sqrt{n}}$ where n is the window length and c some constant. In our data, c varied from $\frac{1}{2}$ to 4, and in most cases we had $c \approx 2$. For the data of Figure 2, the underlying tidal process is certainly not stationary - but the influence of the moon was reduced by calculating β over one month, the annual component was minimized by taking only September. We can roughly estimate σ, with deviations still depending on d.

For medical data, errors were obtained from stationary segments, e.g. deep sleep. Another error of order $\frac{d}{n}$ is observed for very large parameters d (Methods 2,3).

The error estimate $\sigma \approx \frac{2}{\sqrt{n}}$ can be used to determine the appropriate window length. To have the 2σ-radius of the 95% confidence interval for $\beta(d)$ smaller than 0.01, we must take $n \geq 160000$. For Figures 2 and 3 we have $n \approx 8000$ which gives a tolerance of ± 0.04. Thus our method works *only* for big data.

A simple exact method checks whether $\beta(d)$ significantly differs from zero, or the β-curves of two sites in Figure 2 significantly differ at some d. When windows do not overlap, the median test is applicable. Since $\beta(20) > 0$ for all 11 pieces in Figure 3b, one concludes that the median of $\beta(20)$ for the underlying process cannot be zero since the p-value for 11 heads in 11 coin tosses is less than 0.001. Thus, except for γ, most values of the functions in Figure 3 differ significantly from zero.

3. Patterns of length 3. Three equidistant values x_t, x_{t+d}, x_{t+2d} without ties can realize six order patterns. 213 denotes the case $x_{t+d} < x_t < x_{t+2d}$.



For each pattern π and $d = 1, 2, .., T/3$ we count the number $n_\pi(d)$ of appearances in the same way as $n_{12}(d)$. In case $\pi = 312$ we count all $t = 1, ..., T - 2d$ with $x_{t+d} < x_{t+2d} < x_t$. Let S be the sum of the six numbers. Patterns with ties are not counted. Next, we compute the relative frequencies $p_\pi(d) = n_\pi(d)/S$. For white noise it is known that all $p_\pi(d)$ are $\frac{1}{6}$ [8].

As autocorrelation type functions, we now define certain sums and differences of the p_π. The function

$$\tau(d) = p_{123}(d) + p_{321}(d) - \frac{1}{3}$$

is called persistence [17]. This function indicates the probability that the sign of $x_{t+d} - x_t$ persists when we go d time steps ahead. The largest possible value of $\tau(d)$ is $\frac{2}{3}$, assumed for monotone time series. The minimal value is $-\frac{1}{3}$. The constant $\frac{1}{3}$ was chosen so that white noise has persistence zero. The letter τ indicates that this is one way to transfer Kendall's tau to an autocorrelation function. Another version was studied in [22].

Similar to the classical autocorrelation function ρ, a period of length L in a signal is indicated by minima of τ at $d = \frac{L}{2}, \frac{3L}{2}, \frac{5L}{2}, ...$ For these d we have $x_t \approx x_{t+2d}$ so that the patterns 123 and 321 are rare. Near to $d = L, 2L, 3L, ...$ the function τ is large, as ρ. For a noisy sequence, however, a local minimum appears directly at $d = L, 2L, ...$, since patterns tend to have equal probability there. The larger the noise, the deeper the bump. In Figure 3b, τ at $d = 12$ and $d = 24$ shows exactly this appearance and proves the existence of a 24 hour rhythm better than ρ.

Beside τ and β we define two other autocorrelation type functions. For convenience, we drop the argument d.

$$\gamma = p_{213} + p_{231} - p_{132} - p_{312}$$

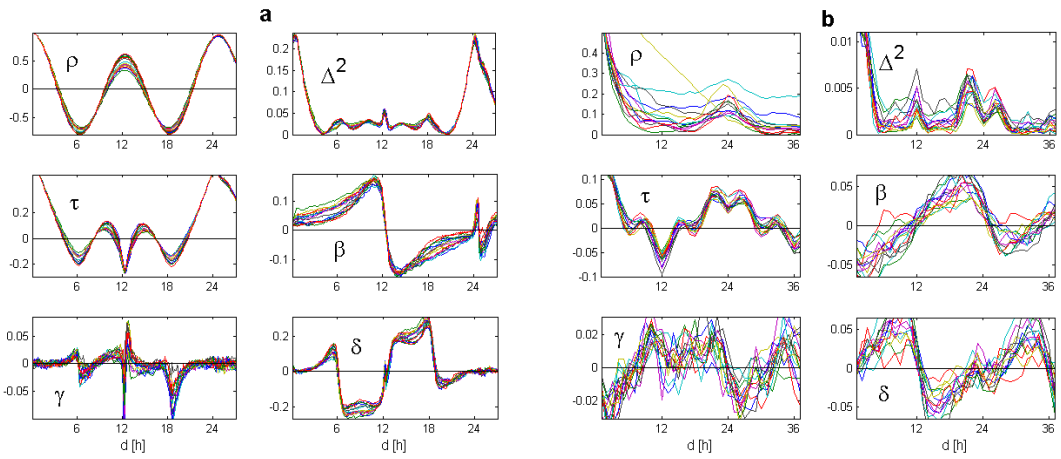


FIGURE 3. Autocorrelation ρ and ordinal functions $\Delta^2, \tau, \beta, \gamma, \delta$ for **a:** the almost periodic series of water levels at Los Angeles (September 1997-2014, data from [20]), **b:** the noisy series of hourly particulate values at nearby San Bernardino 2000-2011 from [21], with weak daily rhythm, see Figure 15. The functions on the left are about three times larger, for Δ^2 nine times, but fluctuations are of the same size.

is a measure of time irreversibility of the process, and

$$\delta = p_{132} + p_{213} - p_{231} - p_{312}$$

describes up-down scaling since it approximately fulfils $\delta(d) = \beta(2d) - \beta(d)$ (Methods 3). Like β, these functions measure certain symmetry breaks in the distribution of the time series. They all have strong invariance properties (cf. Figure 9):

When a nonlinear monotonous transformation is applied to the data – many sensors measure proxy effects which are monotonously, but not linearly related to the target variable – all ordinal functions remain unchanged. They are not influenced by low frequency components with wavelength much larger than d which often appear as artefacts in data. Since τ, β, γ, and δ are defined by assertions like $X_{t+d} - X_t > 0$, they do not require full stationarity of the underlying process. Stationary increments suffice - Brownian motion for instance has $\tau(d) = \frac{1}{6}$ for all d [17]. Finally, ordinal functions are not influenced much by a few outliers. One wrong value, no matter how large, can change $n_\pi(d)$ only by ± 2 while it can completely spoil autocorrelation.

4. Permutation entropy is the Shannon entropy of the distribution of order patterns:

$$H = -\sum_\pi p_\pi \log p_\pi \ .$$

H can be defined for any level n, using the vectors $(x_t, x_{t+d}, ..., x_{t+(n-1)d})$ and their $n!$ order patterns π [8]. In practice we hardly go beyond $n = 7$. Used as a measure of complexity and disorder, H can be calculated for time series of less than thousand values since statistical inaccuracies of the p_π are smoothed out by averaging.

Applications include EEG data [12, 13, 18, 14], optical experiments [2, 1, 11], river flow data [19], control of rotating machines [15, 16], economic applications (cf. [9]), and the theory of dynamical systems [23, 24]. Recent surveys on H are [9, 10]. As a measure of disorder, H assumes its maximum $\log n!$ for white noise. $D = \log n! - H$ is called divergence or Kullback-Leibler distance to the uniform distribution $p_\pi = \frac{1}{n!}$ of white noise.

In this note all functions, including autocorrelation, measure the distance of the data from white noise. For this reason, we take divergence rather than entropy, and we replace $-p_\pi \log p_\pi$ by p_π^2. As before, we drop the argument d. The function

$$\Delta^2 = \sum_\pi (p_\pi - \tfrac{1}{6})^2 = \sum_\pi p_\pi^2 - \tfrac{1}{6}$$

where the sum runs over the six order patterns π of length 3, will be called the *distance of the data from white noise.* Of course, we can define Δ^2 for any length n, replacing $\frac{1}{6}$ by $\frac{1}{n!}$. More precisely, Δ^2 is the squared Euclidean distance between the observed order pattern distribution and the order pattern distribution of white noise. Considering white noise as complete disorder, Δ^2 measures the amount of rule and order in the data. The minimal value 0 is obtained for white noise, and the maximum $\frac{5}{6}$ for monotone time series.

From a practical viewpoint, Δ^2 *is just a rescaling of* H, related to the quadratic Taylor approximation of H around white noise parameters $p_\pi = \frac{1}{6}$

$$H \approx \log 6 - 3\Delta^2 \ .$$

For our data these two functions can hardly be distinguished by eyesight.

5. Partition of the distance to white noise. A Pythagoras type formula combines Δ^2 with the ordinal functions:

$$4\Delta^2 = 3\tau^2 + 2\beta^2 + \gamma^2 + \delta^2 \ .$$

This holds for each $d = 1, 2, ...$ The equation is exact for random processes with stationary increments as well as for cyclic time series. The latter means that we calculate $p_\pi(d)$ from the series $(x_1, x_2, ..., x_T, x_1, x_2, ..., x_{2d})$ where t runs from 1 to T. For real data we go only to $T - 2d$ and have a boundary effect which causes the equation to be only approximately fulfilled (Methods 4). The difference is smaller than 1% in most of our data, see Figures 5, 11, and 16.

This partition is related to orthogonal contrasts in the analysis of variance. When $\Delta^2(d)$ is significantly different from zero, we can define new functions of d:

$$\tilde\tau = \frac{3\tau^2}{4\Delta^2} \ , \quad \tilde\beta = \frac{\beta^2}{2\Delta^2} \ , \quad \tilde\gamma = \frac{\gamma^2}{4\Delta^2} \ , \quad \tilde\delta = \frac{\delta^2}{4\Delta^2} \ .$$

By taking squares, we lose the sign of the values, but we gain a natural scale. $\tilde\tau, \tilde\beta, \tilde\gamma$ and $\tilde\delta$ lie between 0 and $1 = 100\%$, and they sum up to 1. For each d, they describe the percentage of order in the data which is due to the corresponding difference of patterns. Figure 1 shows $\tilde\tau$ for one-minute non-overlapping sliding windows as color code on vertical lines.

For Gaussian and elliptically symmetric processes, the functions β, γ, and δ are all zero, and $\tilde{\tau}$ is 1, for every d [17, Section 5]. An image of $\tilde{\tau}$, as Figure 1, shows to which extent the data come from a Gaussian process and where the symmetry is broken. We get one percentage for each time and delay, 300000 values in Figure 1, and an overall average. As a rule of thumb, we exclude points for which $\Delta^2(d) < \frac{15}{n}$ where n is window length (Methods 2). This bird's view method is a big data counterpart of rigorous tests for elliptical symmetry based on the covariance matrix [25, 26, 27]. For ARMA processes, the average $\tilde{\tau}$ is above 97%, as should be, for EEG data about 80% , for music above 50%, and for heart data only 30-50%. The functions β and δ, and less frequently γ, can represent the main part of Δ^2, but usually only for special values of d. Details are discussed in the supplement.

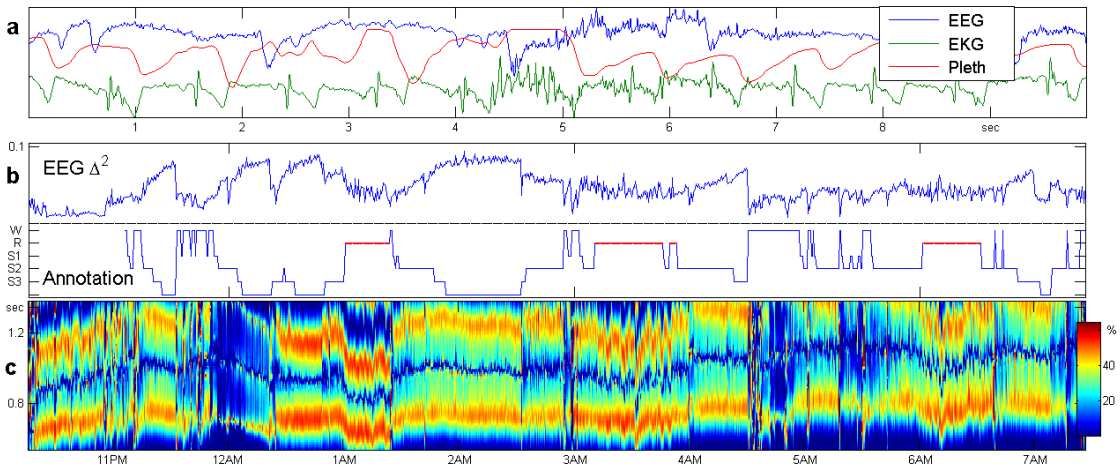


FIGURE 4. Sleep stages of healthy subject n3 from [28]. a: 10 seconds of data from EEG channel Fp2-F4, ECG and plethysmogram. b: Expert annotation of sleep depth from [28] agrees with distance Δ^2. c: Function $\tilde{\beta}$ of the plethysmogram describes changes including REM.

6. Sleep data. As an application, we study data from the CAP sleep database by Terzano et al. [28] available at physionet [3]. Sleep stages S1-S4 and R for REM sleep were already annotated by experts, mainly using the EEG channel Fp2-F4 and the oculogram. Figure 4 demonstrates that permutation entropy Δ^2 of that EEG channel, averaged over $d = 2, ..., 20$ gives an *almost identical estimate of sleep depth.* In [13, 29] permutation entropy was already recommended as indicator of sleep stages. We verified this coincidence in all data of the database with a normal EEG, which seems magic since Δ^2 was introduced as a simple quantity, without any regard to sleep data. For 512 Hz sampling frequency, $d = 2, ..., 20$ corresponds to the time scale between 4 and 40 ms. Thus our Δ^2 is a measure of smoothness of data at small scales which increases when high-frequency oscillations disappear. This seems a complementary viewpoint to the classical R&K rules for sleep stage classification and their recent modification [30] which refer to the appearance of low-frequency waves and patterns.

REM phases are more difficult to detect. The oculogram was used for annotation. Figure 4c presents $\tilde{\beta}$ for the plethysmogram, measured with an optical sensor at the fingertip. We can see all interruptions of sleep and a lot of breakpoints which coincide with changes detected by the annotator. The REM phases are characterized by high values of $\tilde{\beta}(d)$, as well as an increase and a strongly increased variability of the heart rate: the characteristic wavelength d goes down. These variations influence ordinal functions so strongly that changes become more apparent than in a plot of heart rate.

Oximeters measuring the plethysmogram are cheap and easy to apply, in bed at home and possibly in daily life. Many similar devices - armwrists, shirts, smartphone apps - are currently being developed. Our ordinal functions can evaluate and visualize their high-resolution data.

7. Conclusion. On the basis of order patterns, error-resistent autocorrelation functions for big data were introduced. For patterns of length 3, the distance to white noise was divided into four interesting components. Various extensions seem possible. To develop a coherent theory of ordinal time series remains a challenge, despite groundbreaking work by Marc Hallin and co-authors [31, 22, 25] on rank statistics. On the practical side, a spectrogram-like visualization was introduced which allows to see at one glance the course of heart and respiration data over 24 hours (Figures 1,4). Simple as it is, this technique may apply to data of sensors on satellites, weather stations, factory chimneys etc. as well as to experiments on nanoscale [1, 11] which may lay the ground for future generations of sensors. In any application, the methods presented here have to be modified, combined with established techniques, specific machine-learning and imaging tricks. Mathematical extraction of information must keep track with the revolution in sensor and computer technology: this seems an emerging field of study.

References

[1] A. Aragoneses, N. Rubido, J. Tiana-Aisina, M.C. Torrent & C. Masoller, Distinguishing signatures of determinism and stochasticity in spiking complex systems, Scientific Reports 3, Article 1778 (2012)

[2] M.C. Soriano, L. Zunino, O.A. Rosso, I. Fischer & C.R. Mirasso, Time scales of a chaotic semiconductor laser with optical feedback under the lens of a permutation information analysis, IEEE J. Quantum Electronics 47 issue. 2, 252-261 (2011)

[3] A.L. Goldberger et al. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. Circulation 101(23):e215-e220 [Circulation Electronic Pages; http://circ.ahajournals.org/cgi/content/full/101/23/e215]; (2000). Data at: `physionet.org/physiobank/database/capslpdb`

[4] GFZ Seismological Data Archive: `geofon.gfz-potsdam.de/waveform`

[5] P.J. Brockwell & R.A. Davies, Time Series; Theory and Methods, 2nd ed., Springer, New York 1991

[6] R.H. Shumway & D.S. Stoffer, Time Series Analysis and Its Applications, 2nd ed., Springer, New York 2006

[7] A. Davis, M. Rubinstein, N. Wadhwa, G.J. Mysore, F. Durand & W.T. Freeman, The visual microphone: passive recovery of sound from video, SIGGRAPH 2014

[8] C. Bandt & B. Pompe, Permutation entropy: a natural complexity measure for time series, Phys. Rev. Lett. 88, 174102 (2002)
[9] M. Zanin, L. Zunino, O.A. Rosso & D. Papo, Permutation entropy and its main biomedical and econophysics applications: a review, Entropy 14, 1553-1577 (2012)
[10] J. Amigo, K. Keller & J. Kurths (eds.), Recent progress in symbolic dynamics and permutation entropy, Eur. Phys. J. Special Topics 222 (2013)
[11] J.P. Toomey & D.M. Kane, Mapping the dynamical complexity of a semiconductor laser with optical feedback using permutation entropy, Optics Express 22, issue 2, 1713-1725 (2014)
[12] M. Staniek & K. Lehnertz, Symbolic transfer entropy, Phys. Rev. Lett. 100, 158101 (2008)
[13] N. Nicolaou & J. Georgiou, The use of permutation entropy to characterize sleep encephalograms, Clin. EEG Neurosci 2011, 42:24
[14] E. Ferlazzo et al. Permutation entropy of scalp EEG: a tool to investigate epilepsies, Clinical Neurophysiology 125, 13-20 (2014)
[15] U. Nair, B.M. Krishna, V.N.N. Namboothiri & V.P.N. Nampoori, Permutation entropy based real-time chatter detection using audio signal in turning process, Int. J. Adv. Manuf. Technol. 46, 61-68 (2010)
[16] R. Yan, Y. Liu & R.X.Gao, Permutation entropy: a nonlinear statistical measure for status characterization of rotary machines, Mechan. Syst. Signal Processing 29, 474-484 (2012)
[17] C. Bandt & F. Shiha, Order patterns in time series, J. Time Series Analysis 28, 646-665 (2007)
[18] G. Ouyang, C. Dang, D.A. Richards & X. Li, Ordinal pattern based similarity analysis for EEG recordings, Clinical Neurophysiology 121 (2010), 694-703
[19] H. Lange, O.A. Rosso & M. Hauhs, Ordinal pattern and statistical complexity analysis of daily stream flow time series, Eur. Phys. J. Special Topics 222, 535-552 (2013)
[20] The National Water Level Observation Network: `www.tidesandcurrents.noaa.gov/nwlon.html`
[21] California Environmental Protection Agency, Air Resources Board: `www.arb.ca.gov/aqd/aqdcd/aqdcddld.htm`
[22] T.S. Ferguson, C. Genest & M. Hallin, Kendall's tau for serial dependence, Canadian J. Stat. 28, 587-604 (2000)
[23] J. Amigo, Permutation complexity in dynamical systems, Springer 2010
[24] A.M. Unakafov & K. Keller, Conditional entropy of ordinal patterns, Physica D 269, 94-102 (2014)
[25] M. Hallin & D. Paindaveine, Semiparametrically efficient rank-based inference for shape I. Ann. Statistics 34, 2707-2756 (2008)
[26] L. Sakharenko, Testing for ellipsoidal symmetry: a comparison study, Computational Stat. Data Analysis 53, 565-581 (2008)
[27] A. Onatski, M.J. Moreira & M. Hallin, Asymptotic power of sphericity tests for high-dimensional data, Ann. Statistics 41, 1204-1231 (2013)
[28] M.G. Terzano et al., Atlas, rules, and recording techniques for the scoring of cyclic alternating pattern (CAP) in human sleep, Sleep Med. 2 (6), 537-553 (2001).
[29] C.-E. Kuo, S.-F. Liang, Automatic stage scoring of single-channel sleep EEG based on multiscale permutation entropy, Biomedical Circuits and Systems Conference (BioCAS), IEEE 2011, 448-451
[30] C. Iber, S. Anconi-Israel, A. Chesson & S.F. Quan, The AASM manual for the scoring of sleep and associated events: rules, terminologyand technical specifications. Westchester: American Academy of Sleep Medicine (2007)
[31] M. Hallin & M.L. Puri, Aligned rank tests for linear models with autocorrelated error terms, J. Multivariate Analysis 50, 175-237 (1994)

Methods

1. The program code. A few lines of MATLAB code provide all our functions. `x(1:T)` denotes the given time series, `d` is between 1 and dmax.

```
y0=x(1:T-2*d); y1=x(d+1:T-d); y2=x(2*d+1:T);
s=2*((y0>y1)+(y0>y2))+(y1>y2);
```

`s` is a vector which contains the symbols $0, 1, ..., 5$ representing the six order patterns of (x_t, x_{t+d}, x_{t+2d}) for $d = 1, ..., T - 2d$. We correct for missing values NaN and ties by assigning symbols 6 to 11:

```
h=isnan(y0)|isnan(y1)|isnan(y2)|(y0==y1)|(y0==y2)|(y1==y2);
s=s+6*h;
```

Here | means 'or'. If patterns with equality are to be counted, this line has to be modified. We now use the histogram function to determine the frequencies of the six patterns.

```
p=hist(s,12)/(T-2*d-sum(h)); q=p(1:6);
```

This will determine all our functions, for example:

```
beta=p(1)-p(6); tau=p(1)+p(6)-1/3; h2=(q-1/6)*(q-1/6)'; H=-log(q)*q';
```

On a PC, the algorithm obtains Figure 3 within a few seconds and Figure 1 within a few minutes. In some papers [10], permutation entropy of scale d uses patterns for sums of d consecutive terms of the x_t, which makes sense if the x_t represent a density function, like precipitation or workload on a server. This version is easily implemented by cumulative sums, adding `x=cumsum(x)` as first line to the program.

2. Statistical accuracy. We consider a fixed d and windows of length n. Then $p_{12}(d)$ is a random sum of n terms 1,0 which say whether $x_t < x_{t+d}$ is true or false. According to the binomial model, the standard deviation of $p_{12}(d)$ and $\beta(d) = 2p_{12}(d) - 1$ would be $\frac{1}{2}\frac{1}{\sqrt{n}}$ and $\frac{1}{\sqrt{n}}$, respectively. In reality, the variation is bigger because the terms in the sum are correlated. The correlations between differences $Y_t = X_{t+d} - X_t$ are relevant. As a simple model, we take the sum of $\frac{n}{k}$ independent terms 1,0 each of which is repeated k times. For $\beta(d)$ this gives the standard deviation $\frac{c}{\sqrt{n}}$ with $c = k\sqrt{k}$. An estimate of c can only come from data. As a rule we got $c \approx 2$. For τ the variation is a bit smaller, for γ and δ a bit larger. Note that correlations can increase with d. The partition formula and some assumptions give a rough estimate of $\frac{7}{n}$ for the standard deviation of Δ^2. Twice the value, $\frac{15}{n}$, is taken as a kind of confidence bound for the hypothesis $\Delta^2 = 0$. This is just a guideline: for the data in the supplement it works well.

3. Identities for pattern frequencies and boundary effect. We fix d and consider the p_π for the whole time series. The equation $p_{12} = p_{123} + p_{231} + p_{132}$ holds since on the right we count all $t = 1, ..., T - 2d$ with $x_t < x_{t+d}$. Similarly, $p_{12} = p_{123} + p_{213} + p_{312}$ because on the right we count all $t = d + 1, ..., T - d$ with $x_t < x_{t+d}$. Strictly speaking, both relations are not correct since the count for p_{12} is over $t = 1, ..., T - d$. The largest possible difference is $\frac{d}{T-d}$ when we exclude ties for simplicity. This upper bound is negligible for large T and comparably small d, and random fluctuations usually make the difference much smaller than the upper

bound. Taking the difference of the two equations we see that the quantity

$$\epsilon = p_{231} + p_{132} - p_{213} - p_{312}$$

equals zero, with the same degree of accuracy. In other words, the numbers of local maxima and of local minima in a time series coincide. Adding both identities for p_{12} and subtracting $1 = \sum p_\pi$ we obtain

$$\beta = p_{123} - p_{321}$$

which has been used throughout the paper to calculate β, see the code above. Another relation of this type is $p_{12}(2d) = p_{123}(d) + p_{132}(d) + p_{213}(d)$. As a consequence, we get $\beta(2d) = 2p_{12}(2d) - 1 = \beta(d) + \delta(d)$ which says that the functions β and δ are tightly connected. For the mathematical model of processes with stationary increments, as well as for the cyclic time series mentioned in section 5, all these equations are exact [17].

4. Proof of the Δ^2 partition formula. As in the code, we set $q_1 = p_{123} - \frac{1}{6}, q_2 = p_{132} - \frac{1}{6}, q_3 = p_{213} - \frac{1}{6}, q_4 = p_{231} - \frac{1}{6}, q_5 = p_{312} - \frac{1}{6}, q_6 = p_{321} - \frac{1}{6}$. By high-school algebra

$$\begin{aligned} 4\sum q_k^2 = \; & 2(q_1+q_6)^2 + 2(q_1-q_6)^2 \quad + (q_2+q_3+q_4+q_5)^2 + (q_2-q_3-q_4+q_5)^2 \\ & + (q_2+q_3-q_4-q_5)^2 + (q_2-q_3+q_4-q_5)^2 \end{aligned}$$

Since $\sum q_k = 0$, the third square on the right is the same as the first. Using the definitions of $\Delta^2, \tau, \gamma, \delta$ and the identities above for β and ϵ we get

$$4\Delta^2 = 3\tau^2 + 2\beta^2 + \delta^2 + \gamma^2 + \epsilon^2$$

which implies the formula because $\epsilon = 0$. The equation is exact for cyclic time series, and for random processes with stationary increments. For ordinary time series there is a small error, as discussed above, and we must check the accuracy of the equation from the data, see Figures 5, 11, and 16 where differences are below 1% .

5. Contents of supplement. We give numerical and visual evidence for our partition formula and check our bound for small Δ^2. We explore the potential of our methods for various applications - medical data, speech, weather and ecological data - and test how far we can decrease the window size. Autocorrelation and persistence are compared in Figure 9 for a simple AR2-process, in Figure 10 for a song of The Beatles, and in Figure 17 for a laser experiment of Sorriano et al. [2]. Further details of ordinal autocorrelation functions are discussed.

Acknowledgement. I am grateful to Bernd Pompe, Marcus Vollmer, Luciano Zunino, Mathias Bandt, Petra Gummelt and Helena Peña for suggestions and comments on drafts of this paper.

Institute of Mathematics
University of Greifswald
17487 Greifswald, Germany
bandt@uni-greifswald.de

Supplement

1. Overview

We start with a list of our datasets. The essential parameters are window length and mean Δ^2 which quantifies the amount of structure in the data. The latter is not so easy to fix since it depends very much on d. For d near zero, autocorrelation, persistence and hence also Δ^2 assume very large values.

subject / data	window	range of d	mean Δ^2	motivation / application
plethysmogram	3840/ 30s	0.3-1.5s	0.06	proof of partition formula
ECG	15360/30s	0.3-1.5s	0.04	visualize long-term data
EEG	15360/30s	0.25-1.5s	0.002	screening for structure
		0-0.25s	0.013	formula for sleep scoring
AR2 process	2000	1-100	0.01	theoretical model
music/speech	2205/50ms	0-7ms	0.024	speech analysis/synthesis
temperature	720/1month	1-50h	0.03	exploratory data analysis
particulates	1200/50days	1-72h	0.003	notoriously noisy data
tides	1242/5days	10.5-14.5h	0.10	visualization of dynamics
laser data	3520/88ns	741-800	0.01	detection of period

Table 1. Data studied in this supplement

2. Heart and brain data

Figure 5 verifies the partition formula for heart data, and approves our 'confidence bound' $\frac{15}{n}$ for $\Delta^2 = 0$. The vast majority of excluded places comes from artefacts, and it would make no sense to exclude still more places since the partition equation holds true with such a small error. We note that in millions of places in Figure 5 and similar data, $\tilde\tau + \tilde\beta + \tilde\gamma + \tilde\delta$ was always smaller than 100% plus a rounding error of 10^{-10} %. This inequality seems to be generally true.

Processes with small average Δ^2 are more difficult to handle. Figure 6 concerns the EEG channel Fp2-F4 of the same record, n2 of [28], sampled with 512Hz. For $d = 0.25, ..., 1.5s$ the average Δ^2 is only 0.0017, and 47% of the places have small Δ^2. They are spread rather uniformly, so there is little chance to get information from this range of d. We concentrate on $d \le 0.25s$ with average Δ^2 about 0.013 and only 10% places below the bound. There seems some structure connected with sleep annotations. We have chosen the white field near the line $d = 0$ when we average Δ^2 over $d = 4, ..., 40$ms to classify sleep stages. As demonstrated in Figure 7, this rough formula works perfectly well. But Figure 6 indicates other options for a better formula, which certainly can be found by looking more carefully at many data.

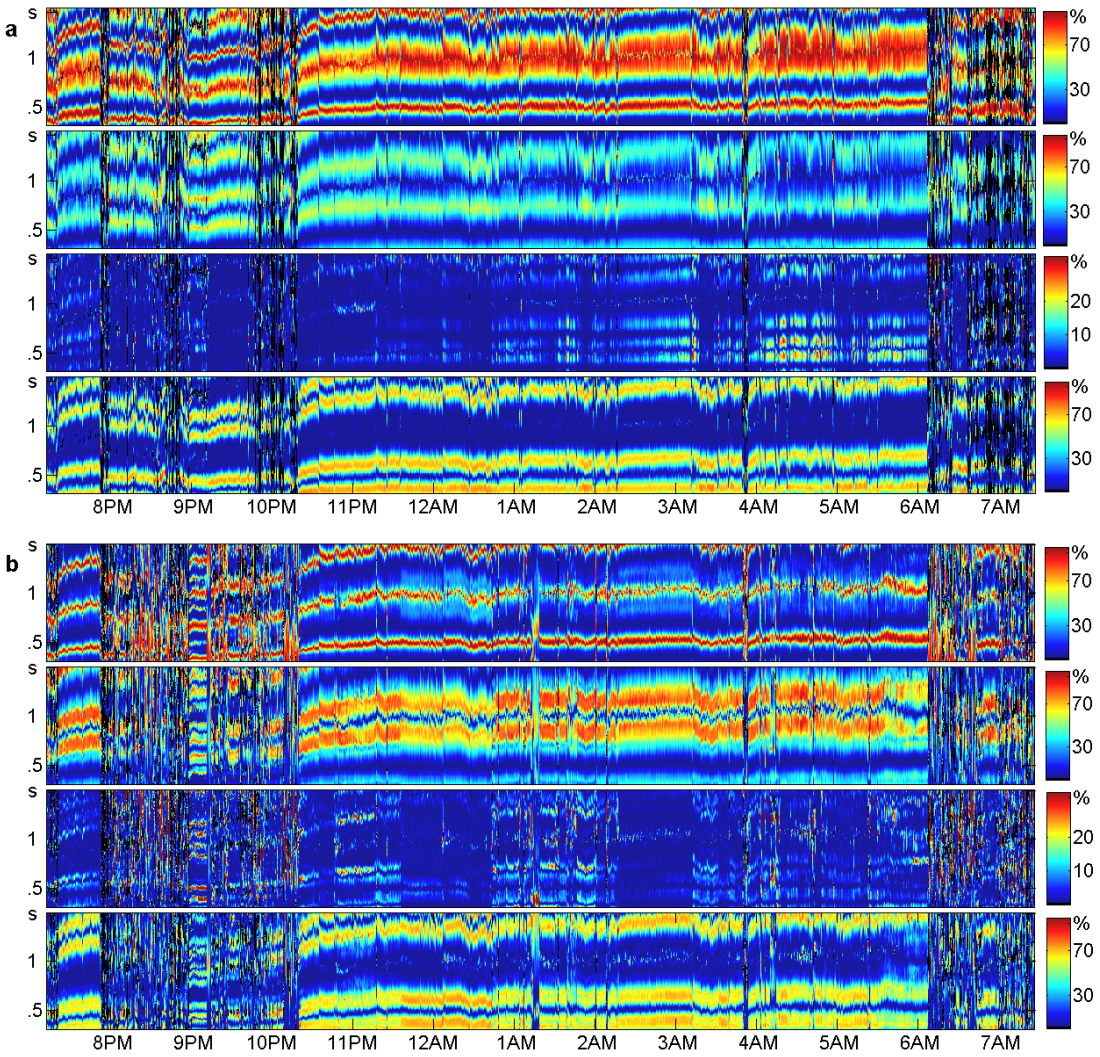


FIGURE 5. Components of Δ^2 for **a** the plethysmogram and **b** the ECG of control n2 of [28], with d from 0.3 to 1.5s. **a** from above: $\tilde{\tau}$ with overall percentage 46% , $\tilde{\beta}$ with 25% , $\tilde{\gamma}$ with 3% (colors scaled by factor 3), and $\tilde{\delta}$ with 26% . The remainder is only 0.07% in average. Places with $\Delta^2 < \frac{15}{n}$ were not included in the average. They account for 7% of all places and are marked black in the picture. The vast majority of black spots is caused by artefacts, movements of the patient. In **b** we have $\tilde{\tau}$ with average 31%, $\tilde{\beta}$ with 35%, $\tilde{\gamma}$ with 5% (colors scaled by factor 3), $\tilde{\delta}$ with 29%, and 0.34% remaining error. The percentage of black spots was 5% . Can you guess REM phases? Solution in Figure 7.

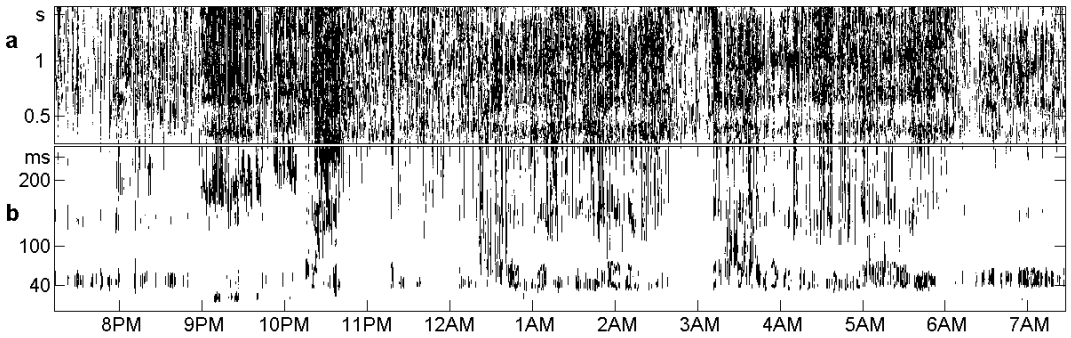


FIGURE 6. Places with $\Delta^2 < \frac{15}{n} \approx 0.001$ in the EEG record of n2 are marked black. **a.** For $d = 0.25...1.5s$ no structure can be seen. **b.** For $d \leq 0.25s$ the light places are related to stages of deep sleep in the figure below.

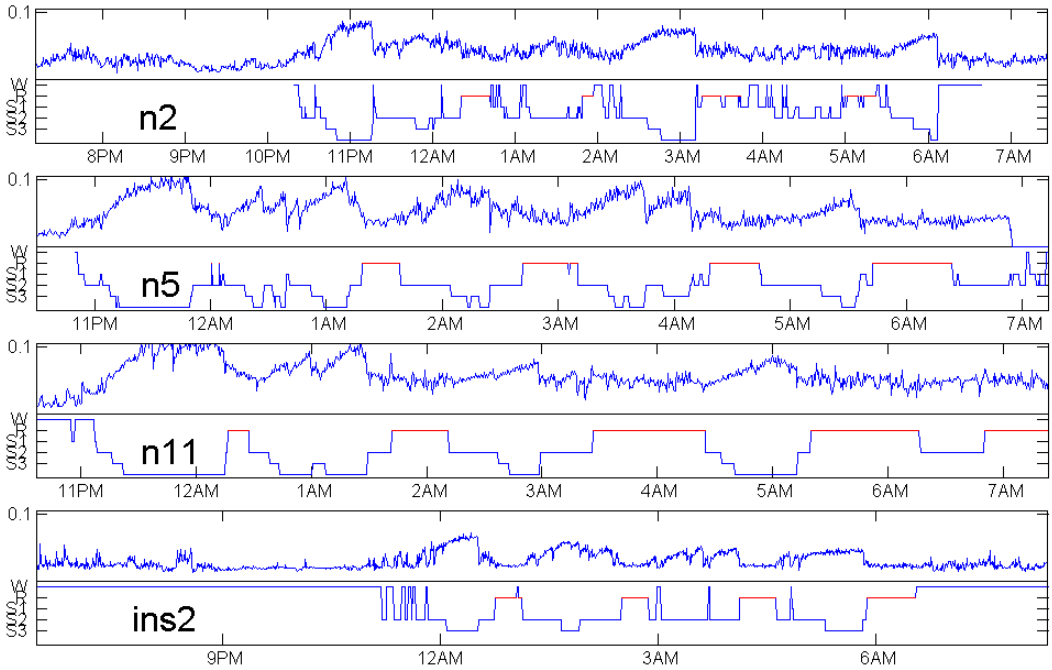


FIGURE 7. Coincidence of mean Δ^2 of EEG and annotation of sleep stages, for three controls from [28] and one patient suffering from insomnia, with smaller Δ^2. Data were selected by record structure, not by the coincidence! There are commercial systems for automatic sleep scoring. But Δ^2 is such a simple, transparent quantity that this coincidence is worthwhile to study. Mean was taken over $d = 2, ..., 20$, that is 4 to 40ms for 512Hz records.

3. A Gaussian model process

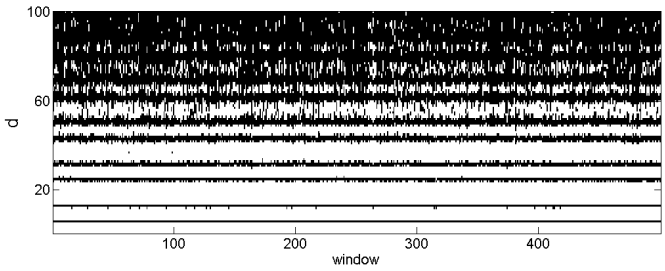


FIGURE 8. Places with $\Delta^2 < 0.0015$ for an AR2 process. For large d they are everywhere, for small d they become concentrated on lines where $\tau(d) \approx 0$.

AR2 processes belong to the simplest stationary Gaussian processes. Autocorrelation and persistence can be determined analytically [6, 17]. Here we use the same procedures as for our data, to find out whether $\tilde{\tau}$ will be really 100%. We take the process $X_t = 1.85X_{t-1} - 0.96X_{t-2} + W_t$ with Gaussian noise W_t which has oscillating and not too rapidly decreasing autocorrelation (Figure 9). Dependencies for $d > 100$ are very small, so we consider d from 1 to 50 or 100. Figure 8 shows that small places for Δ^2 form horizontal stripes where $\tau(d) \approx 0$. We determine $\tilde{\tau}$ on all places and on large places for several thousands of samples of two sizes n.

window	delays	$\Delta^2 < \frac{15}{n}$	$\tilde{\tau}$, all	$\tilde{\tau}$, corrected
$n = 10000$	$d \leq 100$	49%	86%	98.4%
	$d \leq 50$	20%	95%	98.9%
$n = 2000$	$d \leq 100$	76%	73%	97.7%
	$d \leq 50$	55%	87%	97.9%

Table 2. The partition formula is not useful for processes with small Δ^2.

Increasing our bound $\frac{15}{n}$ would only slightly improve $\tilde{\tau}$, on the cost of excluding many other places. Thus our partition formula seems not so useful for small average Δ^2, in this case about 0.01. Here τ is better to work with than $\tilde{\tau}$.

Remark. At this point, let us briefly explain the type of symmetry given when ordinal functions are zero. We take a random process (X_t) with stationary increments, fix a delay d and consider the two-dimensional distribution of $Z_1 = X_{t+d} - X_t$ and $Z_2 = X_{t+2d} - X_{t+d}$. The frequencies of order patterns are probabilities of sectors in the z_1z_2-plane: $\pi_{123}(d) = P(Z_1 > 0, Z_2 > 0), \pi_{321}(d) = P(Z_1 < 0, Z_2 < 0)$ and $\pi_{132}(d) = P(Z_1 + Z_2 > 0, Z_2 < 0), \pi_{213}(d) = P(Z_1 < 0, Z_1 + Z_2 > 0), \pi_{231}(d) = P(Z_1 > 0, Z_1 + Z_2 < 0), \pi_{312}(d) = P(Z_1 + Z_2 < 0, Z_2 > 0)$. When (Z_1, Z_2) has a density function f which is symmetric to the origin, i.e. $f(z_1, z_2) = f(-z_1, -z_2)$ then $\beta(d) = \delta(d) = 0$. If f is symmetric to the line $z_1 + z_2 = 0$, i.e. $f(z_1, z_2) = f(-z_2, -z_1)$ then $\gamma(d) = \beta(d) = \delta(d) = 0$. Ordinal functions measure deviations from this kind of symmetry.

In the figure below, we study how autocorrelation and persistence of the above AR2 process will change when we add various disturbances to the data. While autocorrelation works better when data are perturbed by white noise, persistence performs better for other perturbations which occur in practice. In particular, data from a sensor with nonlinear characteristics need not be calibrated when we use ordinal functions. We note that for cases **c, d,** and **e** in Figure 9, values of $\tilde{\tau}$ from Table 2 remain above 97% .

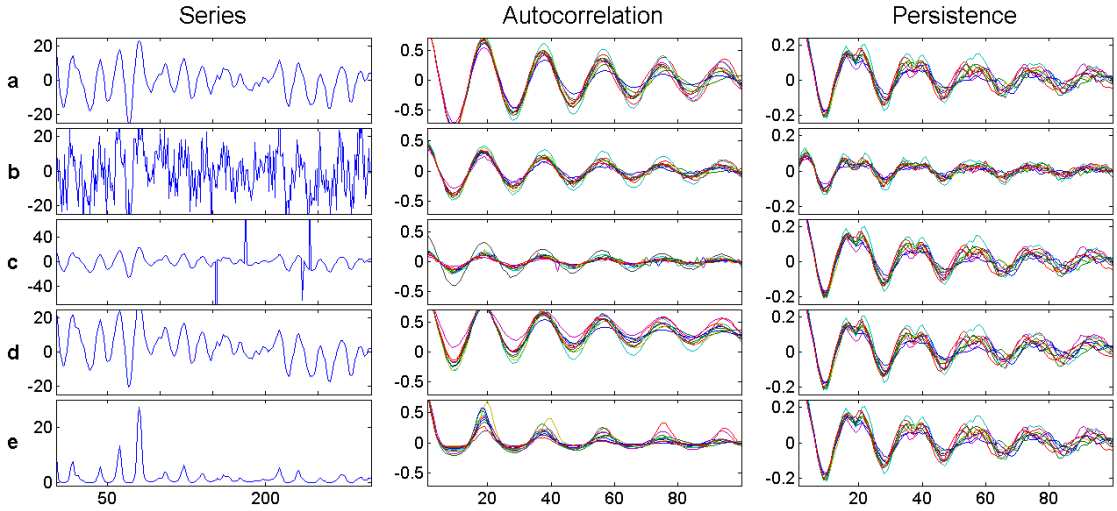


FIGURE 9. Autocorrelation and persistence for an AR2 process with additive disturbances. In each case 10 samples of 2000 points were processed, to determine statistical variation. On the left, 300 points of the respective time series are sketched. **a.** Original signal. **b.** Additive Gaussian white noise with signal-to-noise ratio 1 divides autocorrelation by 2 and makes persistence very flat. **c.** 1% outliers with average amplitude 20 times σ of the original signal will harm the autocorrelation but not persistence. **d.** Low-frequency function $\sin t/300$ added: autocorrelation becomes mostly positive, persistence keeps its shape. **e.** The monotonous transformation $y = e^{x/7}$ is applied to the values: autocorrelation is distorted, persistence unchanged.

4. Speech and music

For music data, the autocorrelogram sometimes shows the melody better than the spectrogram. Persistence behaves similarly.

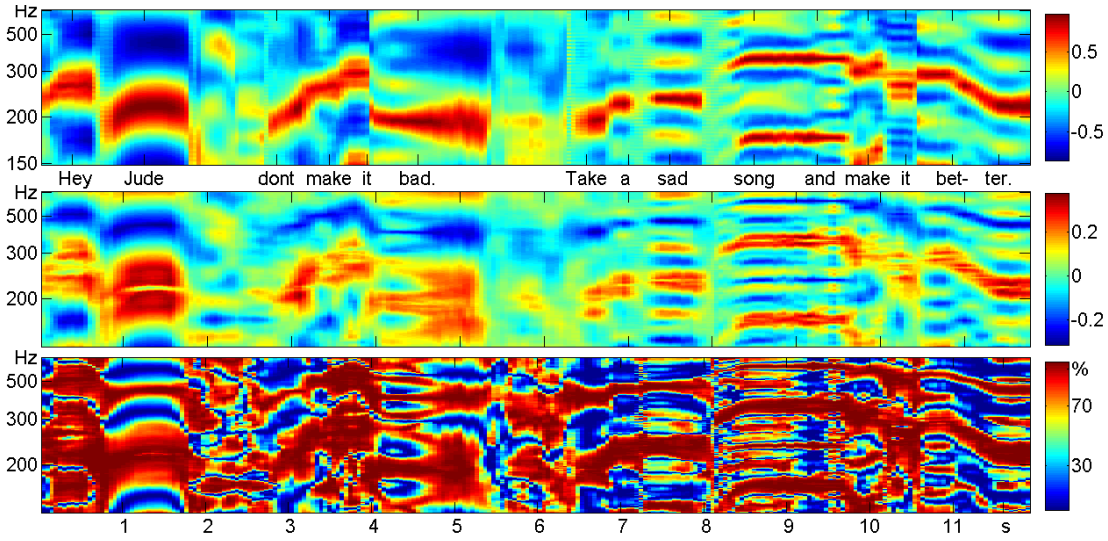


FIGURE 10. The first 12 seconds of the song "Hey Jude" of The Beatles in sliding window analysis. Autocorrelation and persistence were calculated for overlapping windows of 0.25 seconds length, $n = 11025$. Reversed y-axis helps recognize the melody, $d = 49, ..., 294$ corresponds to 900,...,150 Hz. Persistence gives a very similar picture as autocorrelation, except for the local minima at full periods. The function $\tilde{\tau}$ at the bottom has quite a number of blue spots where $\tilde{\beta}$ and $\tilde{\delta}$ yield the main contribution to Δ^2.

If we like to recognize not only melody, but also the text, we have to work with short windows, as shown in Figure 11. The ordinal structure of speech seems an interesting object for its own sake, as well as for applications.

In Figure 12, places with small Δ^2 indicate unvoiced sounds when they are arranged in a vertical pattern. When they form horizontal patterns, which are somewhat wavy in case of a song, they indicate lines $\tau(d) \approx 0$ within voiced sounds. Otherwise they just indicate noise. Voiced sounds can be identified best by their prominent β, δ, and γ-shapes, as shown in Figure 13.

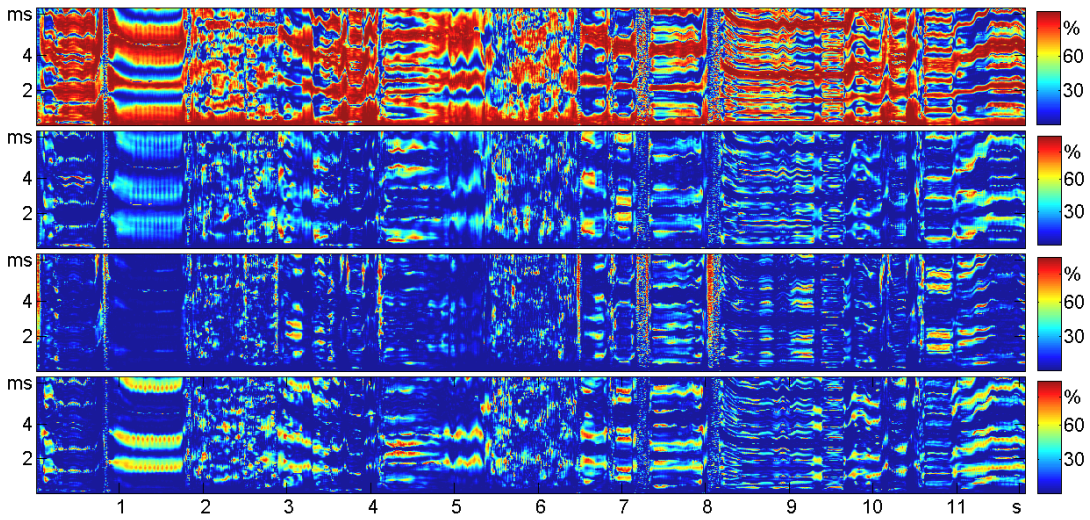


FIGURE 11. The first 12 seconds of "Hey Jude" as partition of Δ^2 with $\tilde\tau, \tilde\beta, \tilde\gamma$, and $\tilde\delta$ (from above). A window of length 50ms was shifted in steps of 10ms. The y-axis is in normal orientation, d runs from 1 to 294, that is, 0.02 to 6.7ms. The short window reveals a lot of detail structure. All vowels show clear β and δ-patterns, even γ is relevant at several places. The overall percentage is 57% for $\tilde\tau$, 15% for $\tilde\beta$, 11% for $\tilde\gamma$, and 16% for $\tilde\delta$. The mean of the rest, i.e. the error of the partition formula, is 0.77% which is good for the short window. Places with small Δ^2 excluded from this calculation are discussed below.

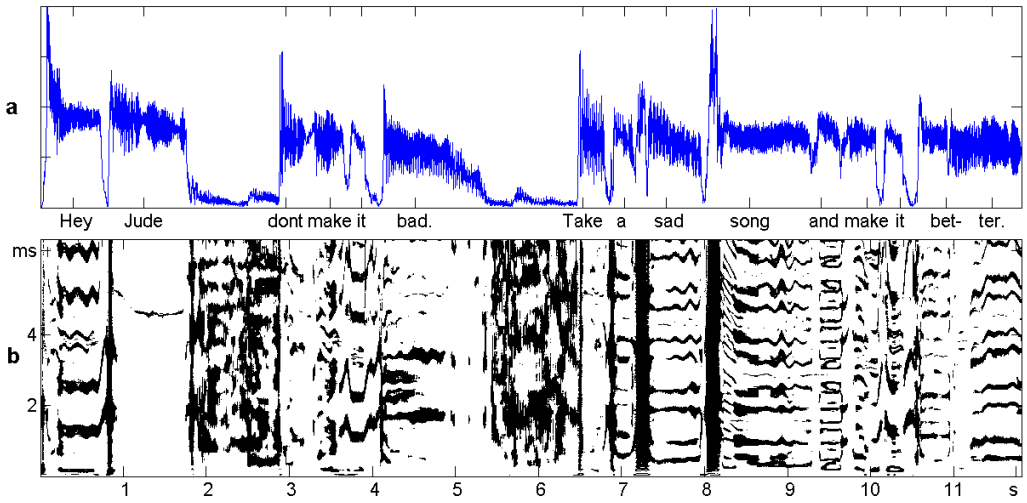


FIGURE 12. In the previous figure, 26% of the places (t, d) fulfil $\Delta^2(d) < \frac{15}{n} \approx 0.007$ so that the signal does not significantly differ from white noise. They form the dark spots in **b.** For comparison, **a** shows the mean absolute amplitude of the signal. We have small Δ^2 for the two breaks with small amplitude, but much more pronounced for the unvoiced sounds, notably 's' in 'sad' and 'song'. In the voiced sounds, places with small Δ^2 appear for d with persistence zero.

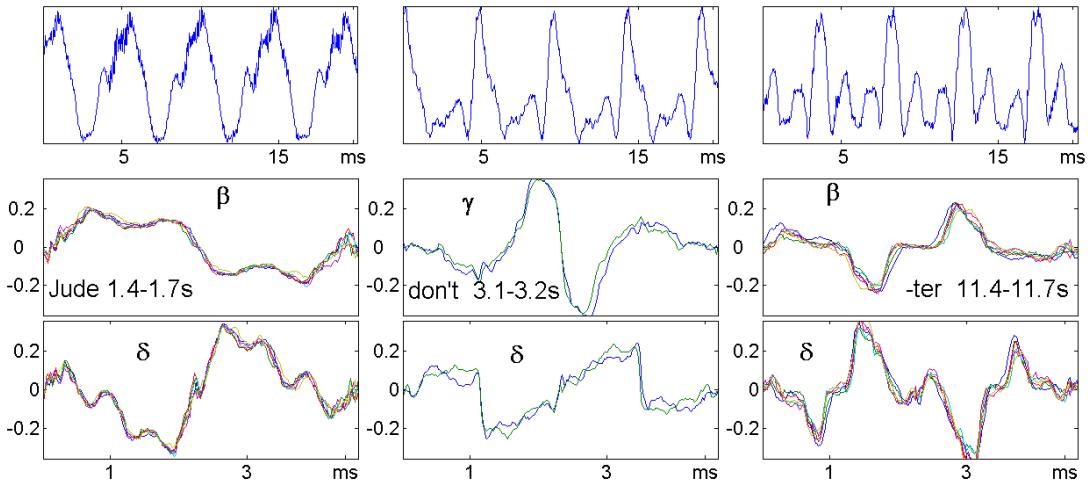


FIGURE 13. Detail from Figure 11. The vowels of 'Jude', 'don't', and the second syllable in 'better' provide stationary parts of the time series lasting for 0.3, 0.1, and 0.3 s. Only 20ms of each signal are shown in the top panel. Ordinal autocorrelation functions were calculated for six resp. two disjoint windows of length 50ms and drawn for one pitch period, which equals 4.5ms for all three sounds.

The study of these patterns can improve methods of speech synthesis, speech analysis, and speaker authentification, but this is not an easy matter. There is a huge variety of β, δ, and γ shapes which have to be collected, classified and understood. And there are open questions. Is there an 'inverse transform' which reconstructs the time series from its ordinal autocorrelation functions? What are the connections between power spectrum, cepstrum and ordinal functions? Can a speaker reliably reproduce the β-shape? Can we hear the ordinal functions? If not, will the computer understand human speech better than humans themselves?

5. Weather, dust, and tides

We conclude with some everyday data. The German Weather Service (`www.dwd.de`, Climate and Environment, Climate Data) provides hourly values of earth temperature, at depth 5cm, starting 1978. Can I learn something about my town? Figure 14 shows the last 2000 values of the year 2013. In autumn, there is still a daily rhythm, in winter this is rarely the case. We determine the persistence of the dataset with sliding windows of length 720, that is, one month. The effect of summer and winter can be seen, and we also see some irregularity: the data were first sampled every 3 years (which one can guess from the picture), then three times a day, and since 2001 every hour. For such plausibility checks even the short window is enough. However, there is no chance to see details in different months. The resolution of data is not sufficient, even for this short window which already causes statistical inaccuracy.

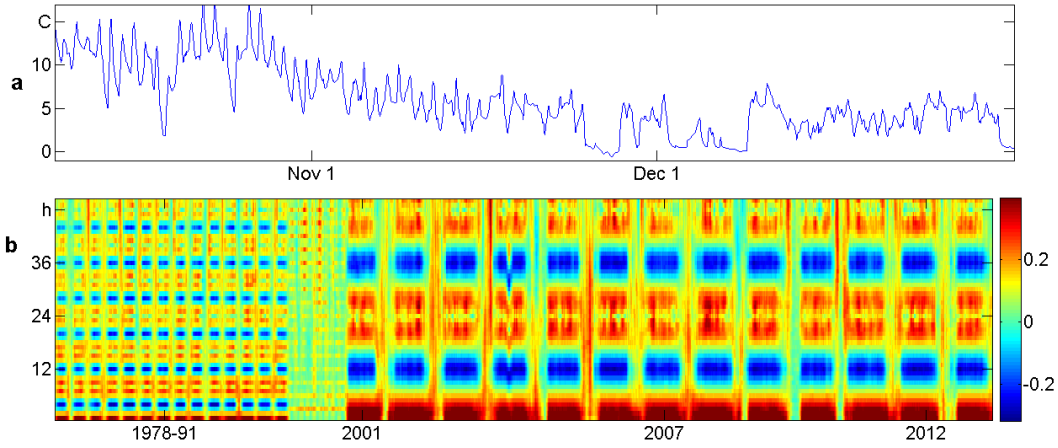


FIGURE 14. Hourly temperature of earth in Greifswald, Germany. **a.** The last 2000 values of 2013 show that the daily cycle is weak in winter. **b.** Persistence for $d = 1, ..., 50$h shows irregularities in the data.

Hourly data of PM10 particulate measurements were already used in Figure 3. These are notoriously noisy data, even if we consider them on logarithmic scale, as seen in the upper part of Figure 15. Particulates in the air are not uniformly distributed, they come in fractal clusters. Actually these data, together with the EEG, have the smallest Δ^2 values in our studies. In the dataset of San Bernardino from [21] (station 3215 Trona-Athol, the other station 3500 shows different characteristics) the mean Δ^2 is 0.003. For our window length 1200, there were 96% of small Δ^2 values according to our bound. Moreover, there were 13% missing values.
For screening the data, we omit all time points with missing values. This is not quite correct, but worth a trial, and the whole financial world works with disrupted time series. Persistence and up-down balance in Figure 15 show that in summer, but not in winter there is a daily rhythm in the data. It is similar as for temperatures above, but not obvious in the data. So the curves in Figure 3b are an effect of summer only.

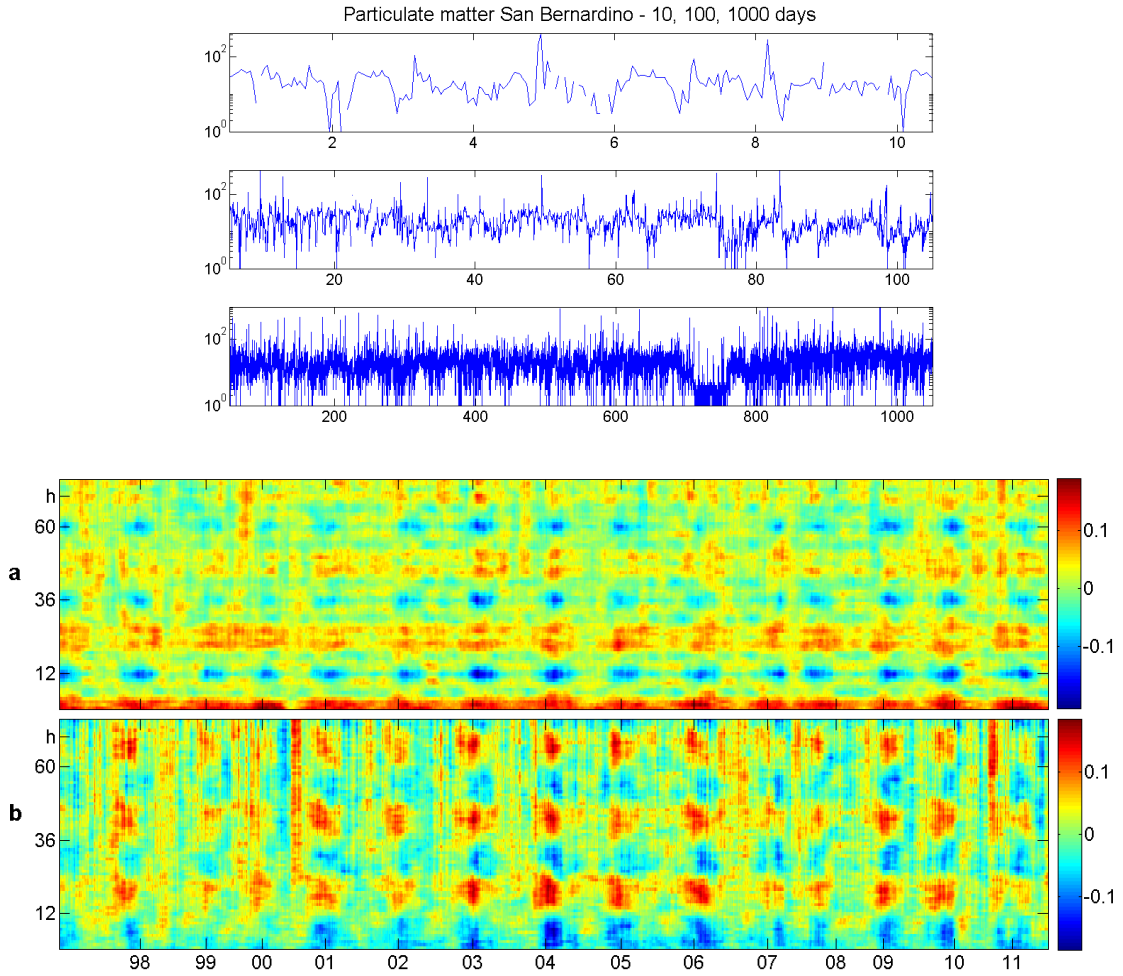


FIGURE 15. Hourly PM10 measurements at San Bernardino, California [21]. The upper panel demonstrates the irregular character of the time series. Ignoring missing data, **a** persistence τ and **b** up-down balance β were determined with a sliding window of length 1200. Years are marked at June 1. There is a daily rhythm in summer which is hardly visible in the data or in autocorrelation, cf. Figure 3. Hourly resolution does not allow to go further into detail.

Based on this impression we can now decide how to proceed the data further. It is not possible, however, to go into detail in Figure 15. With a window of 50 days we can not see changes between weeks. The sensor measurements are done every few seconds, however, and only the hourly averages are kept. With data of finer resolution, say every 6 minutes, much more detail could be studied.

This is demonstrated in Figure 16 for tidal data from Anchorage, providing one more verification for our partition formula. Of course the tides form a very regular process, with the highest Δ^2 of all our examples. The possibility to work with a window of 5 days instead of several weeks, and the resulting details in the figure are due to the high resolution of the data, however. So this note should send a message to colleagues who design measuring equipment: Do not preselect data. Give users an option to study the wealth of all measured values and make their own choices.

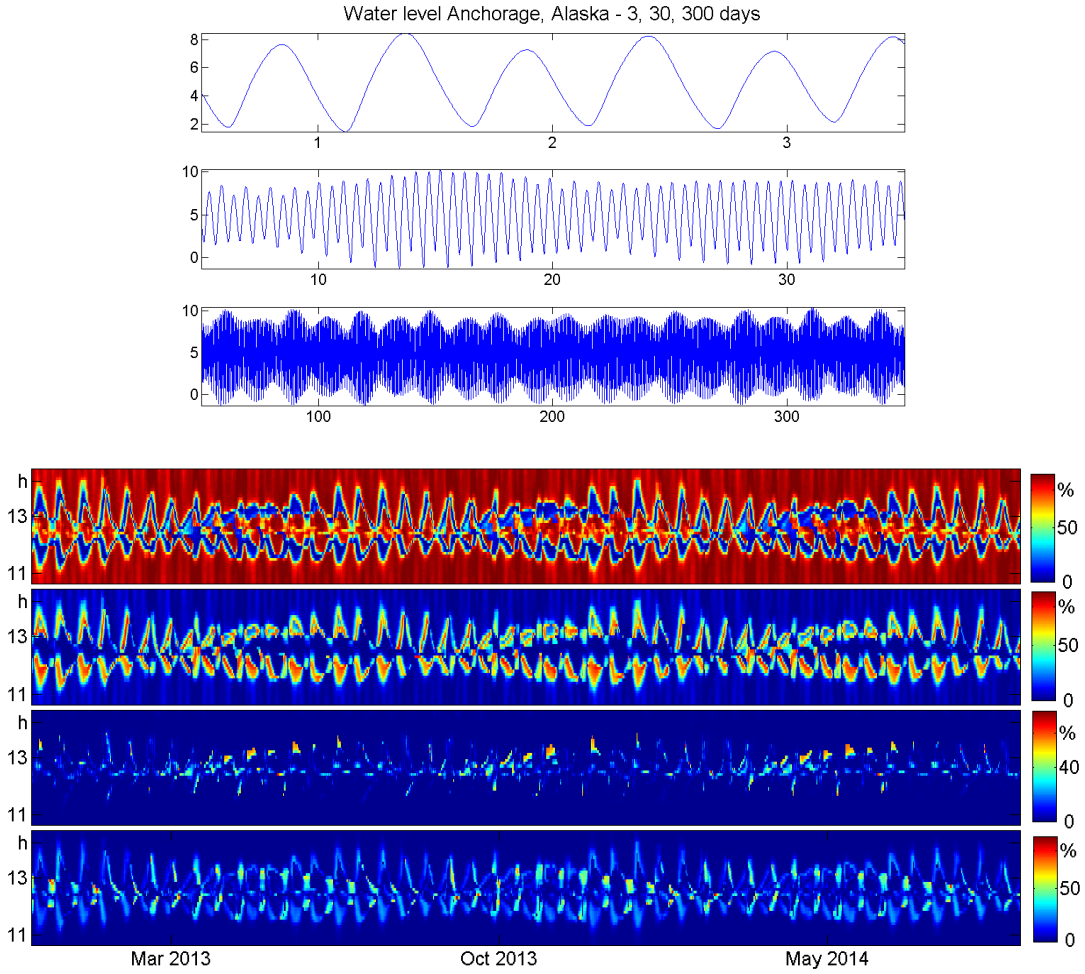


FIGURE 16. We used tidal data to show that there are interesting β-shapes for different sites, by averaging over a whole month. The main actors in the tides' play are moon and sun, however, and the data have a multi-scale structure, shown in the upper panel. It seems interesting that ordinal functions make monthly and yearly cycles more transparent than the series itself, autocorrelation and spectrum. The partition of Δ^2 for water levels from Anchorage was restricted here to the rather short time from Jamuary 2013 to September 2014, and to delays d between 10.5 and 14.5 hours, around half a lunar day. A window of 1242 values, 5 lunar days, was shifted in steps of 12.4 hours. There are only 2% of places with $\Delta^2 < \frac{15}{n}$. If they are excluded, the average value of $\tilde{\tau}$ is 76%, followed by $\tilde{\beta}$ with 15%, $\tilde{\delta}$ with 6%, and $\tilde{\gamma}$ with 2%. The average error is 0.54% .

6. A fast laser experiment

The problem in this last section is rather simple: find the period of a very noisy rhythmic process with extremely short windows. We include this example because the data show what can be measured with current technology within parts of a millisecond.

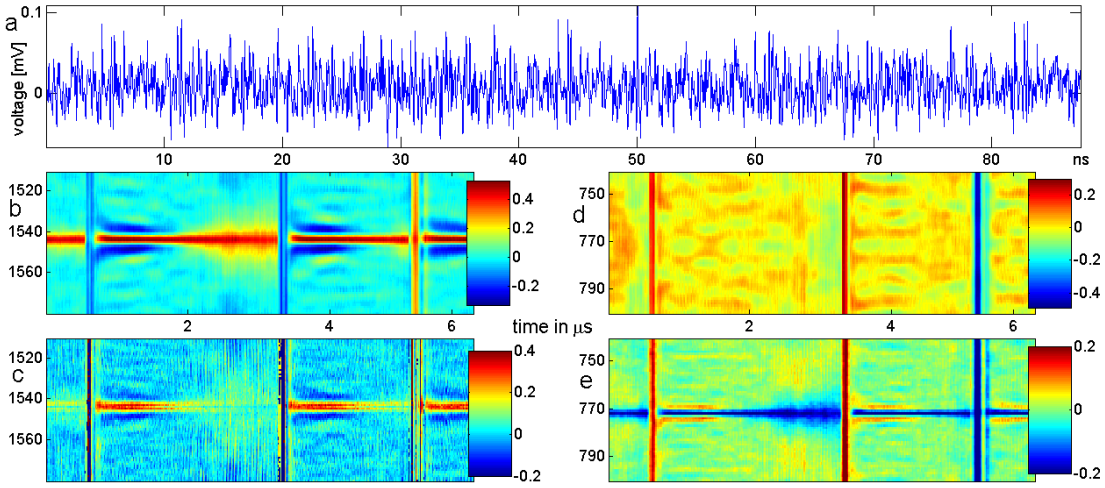


FIGURE 17. Data for this figure come from the laser experiment of Sorriano et al. [2]. A laser with wavelength 1542 nm and 11.7 mA threshold current was excited just above the threshold with 12 mA. The sampling frequency was 25ps which means that 240000 values were measured during $6\mu s$. The optical feedback was arranged so that the signal has a period of $L = 1544$. We try to find this period with a window of 2.3 periods, or 88 ns, just the length of the series shown in **a**, where periodicity can hardly be seen. In **b** and **c** we study autocorrelation and persistence around $d = L$. The vertical interruptions show low-frequency disturbances, the laser does not run smoothly. **d** and **e** show ρ and τ around $d = \frac{L}{2}$. While ρ is a bit better at the full wavelength, τ detects the periodicity reliably at the half wavelength.